

## COLOR IMAGE CHARACTERIZATION FOR PERCEPTION STUDIES

March 2000

Glenn Rogers, Hien Nguyen  
Aberdeen Test Center, APG, MD  
William Pibil  
Army Materiel Systems Analysis Activity, APG MD

### ABSTRACT

This paper proposes a methodology to characterize digital color images for proper reproduction on a monitor. Presenting color images on a monitor in a laboratory reduces the cost of perception studies. Instead of taking subjects to the field for data collection, images collected in the field can be brought to the lab for presentation. This paper details a technique for using data collected in the field to characterize the imaging system and monitor. These data provide a means of accounting for gamma, color balance and resolution for the imaging system/monitor combination.

Gamma, color balance and resolution must be properly managed to render a high quality image. Several approaches for determining gamma response are presented using both field and lab techniques. Accounting for the gamma response of the imaging system and monitor are critical for correct reproduction of contrast. Calibration sources used to determine the gamma of the camera can also be used to color balance the image. If an image is not properly color balanced the shift in hue can confuse the observer. Finally, careful consideration must be given to the effective resolution of the final image. Resolution is probably the most critical factor when studies are being done to extract absolute range information (such as probability of detection vs. range curves). This paper discusses resolution considerations at each stage of the process from real world scene to final reproduction. These considerations include resolution of the sensor, maximum achievable field of view, monitor resolution and distance from the observer to the image on the monitor.

### 1.0 INTRODUCTION

The paper is divided into four sections. The first section is on field data collection. The next section details various methods for measuring the response of the camera using calibration targets. This is followed by a section on color balancing images. The last section covers resolution considerations for both the camera and monitor.

The reference to products or brand names in this paper does not constitute endorsement by the government.

## Form SF298 Citation Data

<b>Report Date</b> <i>("DD MON YYYY")</i> 00032000	<b>Report Type</b> N/A	<b>Dates Covered (from... to)</b> <i>("DD MON YYYY")</i>
<b>Title and Subtitle</b> Color Image Characterization for Perception Studies		<b>Contract or Grant Number</b>
		<b>Program Element Number</b>
<b>Authors</b> Rogers, Glenn; Nguyen, Hien; Pibil, William		<b>Project Number</b>
		<b>Task Number</b>
		<b>Work Unit Number</b>
<b>Performing Organization Name(s) and Address(es)</b> Aberdeen Test Center APG, MD		<b>Performing Organization Number(s)</b>
<b>Sponsoring/Monitoring Agency Name(s) and Address(es)</b>		<b>Monitoring Agency Acronym</b>
		<b>Monitoring Agency Report Number(s)</b>
<b>Distribution/Availability Statement</b> Approved for public release, distribution unlimited		
<b>Supplementary Notes</b>		
<b>Abstract</b>		
<b>Subject Terms</b>		
<b>Document Classification</b> unclassified		<b>Classification of SF298</b> unclassified
<b>Classification of Abstract</b> unclassified		<b>Limitation of Abstract</b> unlimited
<b>Number of Pages</b> 15		

## 2.0 DATA COLLECTION METHODOLOGY

During the image acquisition process all camera settings should be recorded. These include ISO, aperture, exposure time, and focal length settings. This provides information needed to correlate between images, such as multiple images taken to cover a field of regard larger than the field of view of the camera. This also preserves the possibility of characterizing the camera after the image acquisition process has been completed. This can be essential if no color or gray scale references were properly captured in the field.

### 2.1 COLOR BOARD REFERENCES

At a minimum, images of a gray scale board should be taken at various camera settings at some time during the experiment. Ideally, images should be taken every time the camera settings are changed. For photographic film at least one image of a reference should be captured on each roll of film. If possible, the references should be placed along a line of sight between the camera and the object being imaged. This helps to ensure they are characterized under similar front optic illumination conditions as the scene. This is very important for contrast calculations. If there is glare in the front optics it can wash out the scene and reduce calculated contrast based on imager response. However, by having a set of gray scale references in the scene the correct contrast values can be determined.

It is also desirable to include color references on the calibration board. When the images are taken, some of the images may need to be underexposed to assure that in at least one of the images none of the channels, red, green or blue, are clipped.



**Figure 1. Image with red channel clipped on yellow references**



**Figure 2. Underexposed image with no color channels clipped**

In Figure 1 it is not obvious that the red channel is clipped on the yellow reference. Figure 2, which is underexposed by traditional photographic standards, provides an image where none of the channels are clipped on the color standards. Once any channel is clipped, that reference is no longer useful for characterization.

The color board used in this example is a custom configuration obtained from Labsphere. The color coordinates of the colors selected were picked to be c


lose to common outdoor scenes. The white standard, gray scale set and color references are all made from Spectralon® material which is highly lambertian. These are ideal for outdoor use as the acceptable angles (relative to perpendicular) of illumination and measurement are very broad. In addition, these panels can be cleaned and reused many times. The method of obtaining calibrated data to characterize the panels includes the use of a spectroradiometer.

The reference to products or brand names in this paper does not constitute endorsement by the government.


## 2.2 CHARACTERIZING COLOR BOARDS WITH A SPECTRORADIOMETER

The panels are first measured one at a time on a clear sunny day using a spectroradiometer. The color board should be directly illuminated by the sun on a day that has very little variation in illumination during the period the measurements are made. The board should be tilted towards the sky at about a 45° angle and raised up away from the ground. If the board is close to the ground a black cloth should be laid on the ground immediately around the board to prevent excessive ground reflections from entering into the measurements. The spectral power (Watts / meter<sup>2</sup>-steradian) reflected from the panels is plotted against wavelength (nanometers) as shown in Figure 3. The spectral reflectance of the panels can then be calculated by dividing the spectral power of the color panel by that of the white. An example the process is demonstrated using the red panel and the result is shown in Figure 4.

Spectral reflectance curves for all colors and gray scales should be measured in this manner. At this point, all that must be done to provide an absolute color reference input to an image is to take a spectral measurement of the white reference at the same time the image is collected. The spectroradiometric data from the white reference can then be multiplied by the spectral reflectance curves from the color or gray patches to arrive back at spectroradiometric data that would have been measured from the references when the images was taken.



**Figure 3. Spectral power**



**Figure 4. Spectral reflectance**

## 3.0 DETERMINING GAMMA

### 3.1 FIELD CHARACTERIZATION OF IMAGER/SYSTEM RESPONSE

Understanding the gain of the imaging and display systems is critical to the faithful reproduction of contrast. Information derived from the gray scale references (labeled in Figure 5) can be used to map the gain response of the camera as well as the camera/monitor combination. Note: gray scales 6 and 7 were inadvertently placed out of order.



**Figure 5. Gray scale references**


The luminance levels that were measured on the gray scale references in the field, the digital levels of the red, green and blue channels of the image and the output luminance of the monitor are given in Table 1. The responses for the camera alone and the camera/monitor as a system are presented in Figure 6 and Figure 7

The reference to products or brand names in this paper does not constitute endorsement by the government.


respectively. The Kodak camera serial number was DCS 420-4308. The settings at the time this image was captured were ISO: 100, Aperture: F8 and Shutter: 125.

**Table 1. DCS 420-4308 Camera and Monitor Response Data.**

Run	Time	Field (cd/m <sup>2</sup> )	Red	Green	Blue	Monitor (cd/m <sup>2</sup> )
G1	14:05:17	18440	255	255	255	114
G2	14:06:00	11974	255	232	255	101
G3	14:06:32	7861	255	196	221	83
G4	14:07:04	5836	248	181	201	77
G5	14:07:36	4313	222	161	184	61
G6	14:08:12	2304	155	109	126	33
G7	14:09:10	1198	108	73	88	18
G8	14:11:02	440	62	36	52	7




**Figure 6. Camera color channel response**



**Figure 7. System (camera-monitor) response**

Figure 6 shows that the camera response is not linear. Neither is the response of the camera/monitor combination as Figure 7 shows. An additional plot was done to look at the response of the monitor by itself, see Figure 8.

The monitor used was a Gateway 2000, Vivitron 17" monitor with a Matrox MGA Millenium Power Desk PCI video card and associated driver (dated: 8-24-1996). No profile was selected in Windows. The brightness setting was at approximately 50% and the contrast at nearly 100%. Room lights were turned off. The curves for the color channels were generated with a six-step color scale. The gray curve was generated by setting the three channels to the same level.




**Figure 8. Monitor response**

### 3.2 ALTERNATE FIELD CHARACTERIZATION OF CAMERA GAMMA RESONSE

The response curves generated in Figure 6 for the DCS 420 were not smooth. This could be due to the variations in the optics, such as roll off, or the color balance setting or non-uniform illumination across the gray level standards. In order to eliminate the possible effect of spatial variations, a second method was used to try to come up with better curves.

In this experiment, a DCS 420 camera was configured with a 28 mm focal length lens. A series of images were taken of a Spectralon® white board that filled most of the field of view (Figure 9).



**Figure 9. Images of white Spectralon®**



**Figure10. Outline of pixels used on board**



The series of images were taken under clear sunlight conditions with the white board facing the sun and angled to split the difference between the line of sight of the camera and the illumination from the sun. The set of 6 images were taken at varying shutter speeds within a two minute time period. The illumination conditions did not change significantly during this period. Gray levels for the three color channels were obtained by outlining most of the area covered by the white board as shown in Figure10.

The data (taken from Kodak camera DCS 420-1828) are presented in Table 2 and the curves are plotted in linear space in Figure 11. Note that in Figure 11 the curves appear fairly well behaved, i.e. follow a similar pattern (compare to Figure 6). The color balance was set to daylight for the images used to generate Figure 11. These data can also be plotted in logarithmic space to calculate the gamma response of the camera for each color channel. The gamma response is the slope of the linear portion of the curves in log space, see Figure 12. The curves are not quite straight lines over the whole dynamic range in log space. This could be the nature of the response of the camera or, more likley, because the camera shutter speeds are not that precisely controlled.

The reference to products or brand names in this paper does not constitute endorsement by the government.

**Table 2. DCS 420-1828 Camera Response Levels**

Time	Shutter speed (ns)	Red	Green	Blue
13:21:31	16700	255	244	255
13:22:00	8000	232	215	228
13:22:23	4000	185	171	185
13:22:37	2000	120	110	120
13:22:57	1000	81	73	83
13:23:17	500	54	48	59

**Figure 11 Camera response in linear space****Figure 12. Camera response in log space**

Once an adequate gamma response is calculated, the inverse gamma can be found and applied to the data in linear space to “linearize” the response of the system. This should be done before mathematical manipulation of the data such as using Photoshop’s autolevel feature. The gamma and inverse gamma values, calculated from the lower 4 data points of each curve in Figure 12, are presented in Table 3.

**Table 3. Gamma and Inverse Gamma Values**

	Red	Green	Blue
Gamma	0.54	0.56	0.51
Inverse Gamma	1.85	1.80	1.98

### 3.3 LABORATORY CHARACTERIZATION OF CAMERA GAMMA RESPONSE



A laboratory characterization of the camera was done to obtain more data points, under controlled conditions, with the camera settings held constant. This technique included the use of an integrating sphere. The integrating sphere was set to a constant color temperature and the light level was then varied by closing down the port between the light source and the sphere. Images were collected in a light tunnel inside a dark room. The monitoring device that registered the luminance of the sphere was captured in the same image so that a record of the light level became part of the image itself. The gray levels of the color channels are listed in Table 4. The camera used to collect the images was DCS 420-6620. The data are plotted in linear space in Figure 13.

The reference to products or brand names in this paper does not constitute endorsement by the government.



**Table 4. DCS 420-6620 Camera Response From Integrating Sphere**

Luminance	red	green	blue
15.03		251	209
14.02		249	208
13.06		246	205
12.03		243	202
11.05		239	199
10.00		235	195
9.09	251	230	190
8.05	245	225	186
7.09	237	217	180
6.09	228	209	172
5.06	217	199	165
4.02	203	185	154
3.09	173	158	133
2.01	135	122	105
1.02	88	79	73
0.51	57	50	51
0.26	37	31	39
0.11	23	19	27
0	4.41	3	13

**Figure 13. Response in linear space****Figure 14. Response in log space**

These data were also plotted in log space (Figure 14) to look at the gamma response of the camera. It is interesting to note that there seem to be two distinct slopes to the curves in log space. The slopes for the lower portion of the curve, gamma 1 in Table 5, are based on the six data points below the break point in the curve (log luminance of approximately 0.6). For consistency, the slopes for the upper portion of the curves, gamma 2, are based on the first five data points above the break point. The slopes were calculated using a least squares linear regression in Excel. It is also interesting that the gamma of the blue channel is different than the other two channels. This can be seen graphically in Figure 14 and in the data in Table 5. This difference in gain for the blue channel could cause some error when trying to create a systematic mapping for original colors in the scene to colors on a monitor.

**Table 5. Gamma and Inverse Gamma Values**

	Red	Green	Blue
Gamma 1	0.61	0.64	0.48
Gamma 2	0.25	0.25	0.25
Inverse Gamma 1	1.63	1.55	2.09
Inverse Gamma 2	3.98	3.98	4.02

#### 4.0 COLOR BALANCE

##### 4.1 COLOR BALANCING IMAGES

The color balance of the image can be applied using the software that comes with a digital camera or through the use of a gray reference. TWAIN drivers used to import DCS 420 images, for example, allow the user to select various types of lighting conditions such as daylight, tungsten or fluorescent. This method is based on assumptions about the spectral content of the light source illuminating the scene, the response of the camera and the transmission of the optics. This process may work fairly well in some cases but often results in a hue shift that is noticeable. A weakness of this method is that it does not account for the spectral transmission of various lenses. The color balance is applied the same way even if the camera lens or filter is altered. Large discrepancies have been observed between two images taken on the same day under the same lighting conditions using different optics.


A simple technique used to characterize the image, that will account for lens transmission, is to place a gray object in the scene and then to use software to force the red, green and blue channels to be equal on that object. The TWAIN drivers also provide for this type of color balance. Examples of images collected using the daylight illumination assumption versus a gray reference, are presented in Figure 15 and Figure 16. Response curves for the three color channels are plotted just below the images in Figure 17 and Figure 18.




**Figure 15. Daylight color balance**



**Figure 16. Gray reference color balance**



**Figure 17. Daylight color balance**



**Figure 18. Gray reference color balance**

Figure 15 appears too red on a properly calibrated monitor. Looking at the data in Figure 17 (taken from a set of gray scale references), it is obvious that the response of the red channel is too high relative to the green and blue. The image that was color balanced with an external gray reference (Figure 18) looks more like the original scene on a properly adjusted monitor. The data points in Figure 18 are virtually on top of each other for the three color channels over most of the range of luminance levels. This is what should be expected from a properly color balanced image. Simply stated, the red green and blue responses for a color balanced image should be equal on gray objects.

## 4.2 COLOR BALANCING MONITORS


At Aberdeen Test Center an AppleVision 750, 17" monitor is used for perception studies. This monitor was profiled before leaving the factory. It is calibrated weekly using internal sensors and software supplied with the monitor. The AppleVision monitor works in conjunction with Apple's ColorSync® software. According to Apple Computers, ColorSync® is the industry-standard tool for managing color across input, display, and output devices. This system-level software developed by Apple Computer works with scanners, digital cameras, monitors, printers, copiers, proofers, and presses.<sup>1</sup>

To take maximum advantage of this system, the input device should be profiled as well. This can be done using IT8 color targets such as the one shown in Figure 19<sup>2</sup>. Detailed information on profiling cameras

<sup>1</sup> What is ColorSync®, from Apple's website at <http://www.apple.com/coloursync/benefits/>.

<sup>2</sup> From Apple's website at <http://www.apple.com/coloursync/workflows/it8.html>.

and monitors can be found on Apple's web site in a paper titled ColorSync® White Paper. It is located at the web address <http://www.apple.com/colorsync/pdf/colorsyncwhitepaper.pdf>.



**Figure 19. IT8 Color Scanner Target**

#### 4.3 VALIDATING COLOR BALANCE

While the instrumentation is on site, measurements of the targets and background should be collected using a spectroradiometer in conjunction with the color camera. This provides for a validation data set that can be used to serve as a “sanity check” for color mapping techniques used to map colors from the outdoor scene through the camera/monitor system. An example of how this could be done is shown in Figure 20. In this image the areas that were measured are marked and numbered. The 714 spectroradiometer was configured with a one eighth degree field of view to make these measurements. The circles shown were the approximate areas of coverage of the field of view. The spectroradiometric data were taken from the same distance as the camera. This is done to measure the effective color which includes any desaturation or other atmospheric effects.



**Figure 20. Targets in typical background**

The spectroradiometric data taken in the field can be used to calculate “real world” color and luminance information. When the images are displayed on a monitor in the perception lab, the same areas can be measured again using a spectroradiometer or colorimeter to determine color coordinates and luminance information. The spectral curves measured in the field will almost never match those measured off the monitor due to the spiky response of the monitor phosphors. However, if the information is mapped correctly, the resulting color coordinates and relative luminance information should correlate well. If the mapping process is robust it should work over a wide range of camera settings. When making these types of measurements in the field, care should be taken to collect the data under stable clear sunlight conditions. Variations in illumination between the time the image and spectroradiometric data are collected will introduce errors. In the lab, all lights should be turned off in the room to eliminate reflections off the monitor.

The reference to products or brand names in this paper does not constitute endorsement by the government.

## 5.0 DISPLAY RESOLUTION

The primary display resolution consideration is what is the smallest element that can be displayed on the display screen. The size of this element will then determine how far back from the screen the observer has to be to insure the angle extended by the display resolution is smaller than the observers eye resolution capability.

The most common computer display used for laboratory perception experiments are based around the CRT color tube. This device employs three electron guns, a metal screen (shadow mask) near the face of the tube and a matrix of green, blue and red phosphor dots on the CRT face. The shadow mask has a matrix of holes and each of the three guns in sequence are focused at each of the holes. Because the gun beams come in at different angles they strike a different color dot behind the mask. The triad of the blue, green and red phosphor dots (in both directions) are the primary constraint determining the smallest resolution element that can be displayed. The pitch between these dots in millimeters is one of the specifications of the displays. Typical computer displays will have phosphor pitches of 0.28 mm with the latest displays getting down to 0.22 mm. Since the smallest element has to have contributions of red green and blue, this is the smallest element that can be presented to the observer.

The beam of the electron gun is approximately Gaussian in contour and because of small imperfections in aim the beam is larger than the phosphor dots. The result is that the smallest resolution element will be 1.2 times the mask hole pitch (Holst<sup>3</sup>, page 183). As an example, for a 0.28 mm pitch, the smallest possible resolution element (pixel) would be 1.2 X 0.28 or 0.336 mm.

This would only be achieved if the other components of the display such as electronic bandwidth have been optimized to provide the best resolution and could be greater for lower performance displays. If we make the assumption that the human observer capability is 50 cycles per degree for the highest resolution, and that two pixels would be required for a cycle, the next step would be to determine the distance the observer would have to be from the display for the display capability to exceed the eyes capability.

For 50 cycles per degree, one cycle would be 0.02 degrees or 0.348 milliradians. On the display screen would be two pixels or 2 X 0.336 mm equal to 0.672 mm. Using the small angle approximation, the observer would have to be approximately 1.9 meters from the display. Some references have used the approximation of one arc minute (0.29 mrad) as the resolution requirement (Holst, 185) and this would equate to the observer having to be approximately 1.12 meters from the display. This distance would decrease as the masking plate pitch decreases and for a 0.22 pitch, the observer distance would decrease from 1.12 meters to 0.88 meters or 1.49 meters for a .28 pitch.


An additional consideration is the pixel setting of the computer display. The display can be set for different formats such as 800 X 600 (VGA), 1024 X 768(XGA) or 1280 X 960 to name a few. The format is determined by the electronic bandwidth of the system but cannot always be translated to the actual pixel resolution of the display. If the electronic display format exceeds the shadow mask pitch, the shadow mask pitch determines the maximum resolution achievable for the display to the observer. If the electronic display format resolution is less than the shadow mask pitch, the display format will be the limiting resolution capability.

### 5.1 MODULATION TRANSFER FUNCTION APPROACH TO RESOLUTION CONSIDERATIONS


The ideal condition for simulation of scenes for visual perception experiments would be for the spatial response of the display to be one (no degrading of the scene) over the spatial response of the eye. In reality the spatial response of the display degrades as the spatial frequency increases. As an example, the response of a nominal color shadow mask (0.26 mm pitch) CRT display is illustrated Figure 21. For the observer located 1.0 meters from the monitor, the MTF in cycles/milliradian is shown in Figure 22.

---

<sup>3</sup> Holst, Gerald C., CCD Arrays, Cameras and Displays, J C D Publishing, March 1998.




**Figure 21. Color Monitor MTF**




**Figure 22. Monitor MTF at 1.0 meters**

For comparison, the human contrast threshold function for display illumination of 50 cd/m<sup>2</sup> is shown in Figure 23. How do we use this information? The first conclusion is that with present color monitors with the observers at a reasonable distance from the monitor, we are not representing the complete resolution that an observer would see in a natural scene. Is this degrading the scene significantly? I don't think that we can answer this question with the knowledge we have at this time. The primary observation is to be aware of the degradation and try to minimize the impact. In addition, the display resolution parameters such as the display size, display brand and model as well as the observer distance from the display should be reported.

The combination of the display MTF and observer CTF provides an indication of the maximum resolution provided by a perception experiment configuration. As indicated in Figure 24, where the MTF and CTF cross is the maximum spatial frequency that can be considered by the experiment. As can be seen with this configuration, the maximum spatial frequency is approximately 35 cycles per degree where the contrast modulation is approximately 10 percent.



**Figure 23. Contrast threshold function**



**Figure 24. Monitor MTF vs. observer CTF**

## 5.2 DIGITAL RECORDING OF THE IMAGE

The image can be recorded by using film cameras and digital cameras. The resolution capabilities of film cameras are a complex combination of film format, film speed and lens. The image can then be digitized using a scanner. The film format (such as 35mm) determines the area of film being used. The film speed is an indication of the size and density of light sensitive crystals (slower film speeds result in higher resolution). The usual resolution limit is the scanner, normally specified in dots per inch.

The scanner measures the Red, Green, and Blue component for each of the resolution elements. The angular resolution is determined by focal length of the lens used to record the image. The angle subtended by the scanner resolution element on the film and the focal length is the maximum resolution element that can be achieved.

Collecting color images in this manner involves the film development process. The chemicals used in the film development and other factors such as the temperature of the chemical bath can affect the quality and color of the developed film. If film is used to capture the images, at least one color target should be captured with each roll of film.

For the digital camera, a CCD chip is used to transform the incoming light to electrical energy and the level is recorded. The Kodak DCS 420 camera has a CCD array with 1524 X 1012 pixels with a 14.0 X 9.3 mm active area. Each pixel element is a square of 0.00919 mm. The array of pixels is divided up into red, green and blue elements. Therefore, the original image is not actually 1524 X 1012 pixels per channel. Kodak stores the original data from the CCD directly taking about 1.6 megabytes of storage space. A three channel 1524 X 1012, uncompressed image, with no header information, would take approximately 4.6 megabytes of storage space. When the image is imported using the TWAIN drivers it is expanded to 1524 X 1012 pixels/channel. In this process, spatial information is extracted from adjacent color channels to fill out the matrix.

## 5.3 SIMULATING VISUAL SCENES IN A LABORATORY PERCEPTION EXPERIMENT

The basic resolution consideration in presenting the perception experiment is to make sure the image elements subtend an angle smaller than the observers eye capability. This would include both the smallest resolution element on the display and the smallest image resolution element represented on the display as captured by the camera. The starting point is the display/observer geometry. To represent the visual acuity of an observer for a nominal display (0.22 to 0.28 mask pitch) the observer would have to be approximately one meter or more away from the display. To represent magnification, the image has to be increased in size on the screen and the resolution of the image has to be sufficient to permit the increase.

The reference to products or brand names in this paper does not constitute endorsement by the government.

## 6.0 SUMMARY

The methodologies presented in this paper can be used as a guide in designing experiments to collect data for visual perception studies. This paper is far from all-inclusive. There are many ways input and output devices can be characterized for proper reproduction. There are also many options for calibration instrumentation and references. It is not the intention of the authors to endorse any particular product, but rather to give specific examples of the types of hardware and software needed for the task. Following the procedures in this paper should provide the minimum set of information needed to characterize color images for perception studies. It is hoped that this paper will serve as a starting point for the community and that it will be revised and improved to meet individual needs.